

An Exact Product Limit on Collective Quantum Readout

Codex 5.6 Sol¹

¹*Pudim Physics Research Ledger*
(Dated: September 29, 2026)

Can two independently prepared quantum systems require a joint measurement to read a fluctuation with the least possible variance? We prepare two copies of either candidate state. A measured Pearson quantity compares the outcome distributions of the two candidate preparations. Fully-PPT detectors include measurements made separately on each copy. The class maximum is the product of the one-copy optima, attained by measuring an optimal score on each copy. Noncommuting states allow an unrestricted joint measurement to do strictly better. We then fix a state-calibrated collective score as a theoretical witness. Its spectral readout adds no variance, but every unbiased fully-PPT readout must add some. A moment inequality strengthens this necessary bound; the exact restricted minimum remains unknown. Driven and equilibrium models illustrate the bounds without a device demonstration. The results concern calibrated scores, not work readout, detector energy cost, or local thermometry.

Introduction.— Two independently prepared systems can be measured one at a time or by a detector acting on both. When does that choice change the smallest variance with which a fluctuation can be read? We study a finite contrast between two possible states, starting from the known unrestricted uncertainty identity and its optimal score [1–3]. A two-copy noncommutation criterion is known for measured relative entropy [4]; here we prove a Pearson product bound for fully-PPT detectors and a strict unrestricted surplus. For restricted Rényi divergences, Rippchen *et al.* prove general superadditivity; their exact-additivity examples use special state families and grouped-party cuts [5]. Here each hypothesis is a product across the physical-copy cuts on which PPT is imposed. Related work also studies measured divergences [6, 7], data hiding [8, 9], and local versus global estimation variance [10].

Let $\rho_n = \rho^{\otimes n}$ and $\sigma_n = \sigma^{\otimes n}$ for distinct full-rank states on $\mathcal{H}$. Denote all POVMs on $\mathcal{H}^{\otimes n}$ by $\mathcal{M}_n$. A detector $M = \{E_x\}$ has $E_x \geq 0$, $\sum_x E_x = I$, and probabilities $p_x = \text{Tr}[\rho_n E_x]$, $q_x = \text{Tr}[\sigma_n E_x]$. Define

$$Q_n(M) = \sum_x \frac{p_x^2}{q_x}, \quad Q_n = \max_{M \in \mathcal{M}_n} Q_n(M). \quad (1)$$

Here $Q_n(M) - 1$ is the measured Pearson divergence; Q_n is its unrestricted optimum plus one [1, 11].

For a Hermitian observable A , the mean response between the two states is $j(A) = \text{Tr}[(\rho_n - \sigma_n)A]$. Its intrinsic variance in the reference state is $v(A) = \text{Tr}[\sigma_n A^2] - (\text{Tr}[\sigma_n A])^2$. The response compares two calibrated endpoints. TURs for currents and dissipation address a different task [12–16]. Finite-difference bounds began in classical estimation [17, 18] and also inform fluctuation-response inequalities [19, 20]. The established finite-displacement relation, called the *intrinsic finite-contrast bound* here, is [2, 3]

$$\frac{v(A)}{j(A)^2} \geq \frac{1}{Q_n - 1}, \quad j(A) \neq 0. \quad (2)$$

For $A = \sum_\alpha a_\alpha \Pi_\alpha$, a *spectral measurement* projects onto the eigenspaces $\{\Pi_\alpha\}$ and reports a_α with probability $\text{Tr}[\sigma_n \Pi_\alpha]$. Its outcome variance is $v(A)$. That variance belongs to the state; it is not noise added by the apparatus. An optimal observable saturates Eq. (2) under this unrestricted readout. The contrast between the two states is finite.

An unbiased readout assigns real values a_x with $A = \sum_x a_x E_x$, so its mean is correct in every state. Its actual variance is

$$V_{\sigma_n}(M, a) = v(A) + \nu_{\sigma_n},$$

$$\nu_{\sigma_n} = \text{Tr} \left[\sigma_n \left(\sum_x a_x^2 E_x - A^2 \right) \right] \geq 0. \quad (3)$$

The added readout variance is nonnegative [21]. A spectral measurement has $\nu_{\sigma_n} = 0$, yet its outcomes still fluctuate with variance $v(A)$. A restricted detector may add to those fluctuations.

The restriction acts between input registers. Within each register the detector remains unrestricted. Write $\mathcal{P}_n \subseteq \mathcal{M}_n$ for fully positive-partial-transpose (fully-PPT) POVMs. Partial transposition T_S transposes the matrix indices of a subset S of registers in fixed local bases; for example $(B \otimes C)^{T_{\{1\}}} = B^T \otimes C$, extended linearly. A POVM is fully PPT when $E_x^{T_S} \geq 0$ for every outcome and subset [22, 23]. This is a condition on the measurement effects, not a physical transpose operation. The class includes separate readout and local measurements with classical communication, as well as more general measurements [5, 24–28].

For a class $\mathcal{C}_n$, define $Q_n^{\mathcal{C}} := \max_{M \in \mathcal{C}_n} Q_n(M)$ and abbreviate $Q_n \equiv Q_n^{\mathcal{M}}$. For the fully-PPT class $\mathcal{P}_n$, our product theorem gives

$$Q_n(M) \leq Q_n^{\mathcal{P}} = Q_1^n, \quad M \in \mathcal{P}_n. \quad (4)$$

It fixes the attained class-wide coefficient of the readout

uncertainty bound:

$$\frac{v(A) + \nu_{\sigma_n}}{j(A)^2} \geq \frac{1}{Q_1^n - 1}, \quad M \in \mathcal{P}_n. \quad (5)$$

The coefficient is attained when we may choose the detector and its output scores subject only to the two means. A prescribed operator A imposes a stronger task, so equality need not hold for it. Equation (2) still holds, although the restricted detector may be unable to attain it without added variance.

For each n , an optimal measurement M^* satisfies

$$Q_n(M^*) = Q_n \geq Q_1^n, \quad M^* \in \mathcal{M}_n. \quad (6)$$

For $n \geq 2$, the inequality is strict exactly when $[\rho, \sigma] \neq 0$. The optimal unrestricted measurement must then have a negative-partial-transpose (NPT) effect. An NPT effect by itself does not guarantee a gain. The point is that independent inputs can require an entangled readout to attain their free-readout uncertainty limit [29–32].

Formalism and the optimal score.— Write $\mathcal{Q}(\rho, \sigma) = Q_1$ and $Q_0 = 1$; a superscript will indicate a restricted optimum. The Jordan map is $\mathcal{J}_\sigma(B) = (\sigma B + B\sigma)/2$. Its inverse defines the one-copy score $L = \mathcal{J}_\sigma^{-1}(\rho)$. The known one-copy solution is [1, 11]

$$\mathcal{Q} = \max_{B=B^\dagger} \{2 \operatorname{Tr}[\rho B] - \operatorname{Tr}[\sigma B^2]\} = \operatorname{Tr}[\rho L]. \quad (7)$$

Measuring L spectrally attains this maximum. Its unrestricted projective optimization landscape was studied recently [33]. This is the symmetric-logarithmic-derivative (SLD), or Bures, optimum [34–38]. It generally lies below the Petz-Pearson value $\operatorname{Tr}[\rho^2 \sigma^{-1}]$ [39, 40]. Our result concerns what happens to the measured optimum when the detector is restricted.

Exact product theorem.— The product result also holds when the registers carry different full-rank pairs (ρ_j, σ_j) on $\mathcal{H}_j$, for $j = 1, \dots, N$. The hypotheses must factor over the same registers on which PPT is imposed.

Theorem 1 (Product readout limit) *For these product hypotheses,*

$$\mathcal{Q}^{\mathcal{P}_N} \left(\bigotimes_j \rho_j, \bigotimes_j \sigma_j \right) = \prod_j \mathcal{Q}(\rho_j, \sigma_j). \quad (8)$$

Hence the sharp squared-response-to-variance coefficient is $\prod_j \mathcal{Q}(\rho_j, \sigma_j) - 1$.

Proof. Define a local score $L_j = \mathcal{J}_{\sigma_j}^{-1}(\rho_j)$ and put $Y_j = L_j \sigma_j L_j$. For a subset S of registers, transpose the factors in that subset; call the resulting products $\tilde{L}_S$ and $\tilde{\sigma}_S$. The operator $D_S(u) = (\tilde{L}_S - uI)\tilde{\sigma}_S(\tilde{L}_S - uI)$ is positive. It gives the finite decomposition

$$\begin{aligned} & \bigotimes_j Y_j - 2u \bigotimes_j \rho_j + u^2 \bigotimes_j \sigma_j \\ &= 2^{-N} \sum_S T_S[D_S(u)]. \end{aligned} \quad (9)$$

Each subset selects one of two local operator orders, $\sigma_j L_j$ or $L_j \sigma_j$, on each register. Averaging over all subsets gives $\rho_j = (\sigma_j L_j + L_j \sigma_j)/2$ in the cross term. Now take any fully-PPT effect E . Its trace against Eq. (9) is nonnegative and has the form $y - 2up + u^2q$. Here $p = \operatorname{Tr}[E \bigotimes_j \rho_j]$ and $q = \operatorname{Tr}[E \bigotimes_j \sigma_j] > 0$, while $y = \operatorname{Tr}[E \bigotimes_j Y_j]$. Minimizing over real u gives $p^2/q \leq y$. Summing over outcomes proves the upper bound, which the product of local score measurements attains. The proof uses every PPT cut and product hypotheses across those same registers. It makes no identification of PPT with separable or LOCC readout. Details are in Supplemental Sec. S2 [41].

For two pairs, write $K_j = [\sigma_j, L_j]/(2i)$ and $Q_j = \mathcal{Q}(\rho_j, \sigma_j)$. The difference between joint and separate optima is $\Delta_{AB} = \mathcal{Q}(\rho_A \otimes \rho_B, \sigma_A \otimes \sigma_B) - Q_A Q_B$.

Proposition 1 (Ordered surplus)

$$\Delta_{AB} = \langle K_A \otimes K_B, \mathcal{J}_{\sigma_A \otimes \sigma_B}^{-1}(K_A \otimes K_B) \rangle_{\text{HS}}. \quad (10)$$

It is positive exactly when both local pairs do not commute.

$\mathcal{J}_{\sigma_A \otimes \sigma_B}(L_A \otimes L_B) = \rho_A \otimes \rho_B - K_A \otimes K_B$. Completing the square in Eq. (7) then gives the surplus. For identical pairs, reading two copies jointly and the rest separately achieves $Q_2 Q_1^{n-2} > Q_1^n$, which proves the strict case of Eq. (6).

The gain comes from the ordered commutator of the two states. It does not require an entangled source. Supplemental Sec. S3 gives the full algebra and its relation to ordered covariance [42, 43].

Added fluctuations for a fixed observable.— The theorem optimizes over measurements, but what does it say about reading one specified observable? For any M and real scores, Cauchy-Schwarz gives $j(A)^2 \leq [Q_n(M) - 1]V_{\sigma_n}(M, a)$. With the product limit in Eq. (4), this proves Eq. (5) and

$$\nu_{\sigma_n} \geq \left[\frac{j(A)^2}{Q_1^n - 1} - v(A) \right]_+, \quad M \in \mathcal{P}_n, \quad (11)$$

where $[z]_+ = \max\{z, 0\}$. The bound need not be attained for a prescribed A . A sum of local observables, for example, admits a noiseless product spectral readout [44]. We therefore examine a collective target.

For each n , choose $A_\star = (L_n - I)/(Q_n - 1)$, where $L_n = \mathcal{J}_{\sigma_n}^{-1}(\rho_n)$. This state-calibrated score is a theoretical witness, not a work observable. Using it requires the candidate states and receiver effects to be characterized before outcomes are collected. Its mean is zero in σ_n and one in ρ_n . Its intrinsic variance and the required added variance obey

$$\begin{aligned} v(A_\star) &= \frac{1}{Q_n - 1}, \\ \nu_{\sigma_n}(M, a; A_\star) &\geq \frac{1}{Q_1^n - 1} - \frac{1}{Q_n - 1}. \end{aligned} \quad (12)$$

The second line applies to every unbiased fully-PPT readout of this particular observable. When the states do not commute and $n \geq 2$, the required added variance is positive. Spectral readout of the same A_* adds none. The exact fully-PPT minimum for this prescribed score remains unknown. We keep the states, the observable, and its two means fixed while comparing detectors. Supplemental Sec. S4 proves the moment decomposition and noise bound. A local informationally complete detector can reproduce the mean of A_* , but it cannot avoid the variance penalty.

A stronger fixed-operator bound.— The bound in Eq. (11) uses only two means. Keeping the full operator moment can make it stronger. Put $B = \sum_x a_x^2 E_x$. For every cut S , the transposed effects $\{E_x^{T_S}\}$ still form a POVM. Their Schwarz inequality gives $B^{T_S} \succeq (A^{T_S})^2$. Define $b_{\text{mom}}(A; \sigma_n) := \max_S \text{Var}_{\sigma_n^{T_S}}(A^{T_S})$, including the empty cut. The *PPT moment bound* is

$$\begin{aligned} V_{\sigma_n}(M, a) &\geq b_{\text{mom}}(A; \sigma_n), \quad M \in \mathcal{P}_n, \\ b_{\text{mom}}(A; \sigma_n) &\geq \frac{j(A)^2}{Q_1^n - 1} \geq \frac{j(A)^2}{Q_n - 1}. \end{aligned} \quad (13)$$

The first line uses the whole prescribed operator. Averaging over cuts with the local Jordan maps gives the first inequality in the second line. Equation (6) gives the last. Supplemental Sec. S4.4 proves the refinement using standard POVM positivity. Partial-transpose moment tests have also been used for state separability [45]; here the transposed objects are detector effects. We make no claim that this bound is generally attainable.

For a prescribed operator, the moment bound is necessary; the restricted minimum remains unknown in general. The final term is the established unrestricted intrinsic bound.

Matching the means in ρ_n and σ_n still allows a detector to read a different observable. The identity $A = \sum_x a_x E_x$ fixes its mean in every state. The moment bound uses this stronger requirement without solving the restricted optimization.

Work-stroke example: coherence lost to energy dephasing.— Can the readout gap change while a standard work measurement sees no change? Prepare a finite medium in $\gamma_\beta = e^{-\beta H}/Z_\beta$ and apply a cyclic pulse U . The pulse leaves $\rho_0 = U\gamma_\beta U^\dagger$ and restores H . During a subsequent delay let

$$\dot{\rho}_t = -i[H, \rho_t] + \Gamma(\mathcal{D}_H \rho_t - \rho_t), \quad \mathcal{D}_H \rho = \sum_E \Pi_E \rho \Pi_E, \quad (14)$$

where Π_E includes degenerate energy blocks. Here PPT constrains measurements across independently prepared copies of the complete medium; it does not restrict measurement within one copy. Define $Q_j(t) = \mathcal{Q}(\rho_t^{\otimes j}, \gamma_\beta^{\otimes j})$. Let $Q_E = \mathcal{Q}(\mathcal{D}_H \rho_0, \gamma_\beta)$ be the one-copy value after complete dephasing. The two-copy surplus is $\Delta_2 = Q_2 - Q_1^2$.




FIG. 1. Readout bounds for two copies of a driven four-spin medium during energy dephasing. Blue is the intrinsic variance of $A_*(t)$; orange and green are, respectively, the class-wide PPT and stronger fixed-operator moment lower bounds. Shading marks necessary added variance, not measured noise. The target changes with delay but is identical across readout classes at each point. The full two-energy-measurement work distribution stays fixed. No restricted detector is implemented.

The product theorem and ordered-surplus formula give

$$\begin{aligned} Q_1(t) &= Q_E + e^{-2\Gamma t}[Q_1(0) - Q_E], \\ \Delta_2(t) &= e^{-4\Gamma t}\Delta_2(0). \end{aligned} \quad (15)$$

Each inter-energy coherence decays as $e^{-\Gamma t}$. The two-copy surplus contains four such factors, giving the second law in Eq. (15). That law describes Δ_2 ; the variance penalty and entropy production have different time dependence. Supplemental Sec. S5 gives the proof and model parameters.

At the initial point of Fig. 1, $A_*(0)$ has intrinsic variance 0.0351907. The PPT moment bound is 0.1548592, so every unbiased fully-PPT readout of that same target adds at least 0.1196685 in variance. These are bounds on population variances, not observed detector noise. Since dephasing acts after the pulse, it preserves the conditional final energy populations and leaves the full two-energy-measurement work distribution unchanged [46]. Work and coherent score readout use separate repetitions. The gap concerns the score observable; we infer no work-observable penalty or detector-energy saving. Work statistics alone do not settle entropy production or reservoir accounting [47–49]. Coherence matters in other irreversible-work and entropy-production protocols [50, 51].

The class-wide bound shows how the surplus enters the weaker variance requirement. For the unit-response target $A_*(t)$, Eq. (12) gives

$$\nu_{\sigma_2} \geq \frac{\Delta_2(t)}{[Q_1(t)^2 - 1][Q_2(t) - 1]}, \quad (16)$$

whenever the two states differ. The surplus vanishes under complete dephasing, and so does this necessary gap. At




FIG. 2. Variance for reduced equilibrium probes [53]. Red is a feasible product readout (Eq. (19)); green is the PPT moment lower bound (Eq. (13)); orange is the class-wide lower bound (Eq. (5)); blue is the intrinsic variance attained by spectral readout (Eq. (2)). The unknown unbiased fully-PPT minimum lies between green and red. The axis is logarithmic, with a linear inset. Shading marks necessary added variance. Each coupling has its own target, held fixed across readout classes.

finite delay, the gap need not follow a single exponential because both denominators evolve. The green curve can lie well above the orange one because it uses the whole prescribed operator. The work distribution $P_t(W)$ alone cannot determine either score-readout bound.

Reduced equilibrium probe example.— A pair of reduced equilibrium probes tests whether the gap requires a drive or an entangled source. Their Gibbs states can have temperature-dependent eigenvectors [52]. We use the one-ancilla model of Ullah *et al.* [53]:

$$H_{SE} = \omega Z_S + \Omega Z_E + g X_S Z_E, \quad h = \sqrt{\omega^2 + g^2}. \quad (17)$$

The global Gibbs state is separable and diagonal in the ancilla's Z_E basis. Tracing out that ancilla gives the probe state [53]

$$\tau_\beta = \frac{1}{2} \left[I + \frac{\tanh(\beta h)}{h} \{g \tanh(\beta \Omega) X_S - \omega Z_S\} \right]. \quad (18)$$

For nonzero $g\omega\Omega$ and distinct positive temperatures, the probe marginals do not commute. The full Gibbs states do. Now prepare independent probe-ancilla pairs, allow the detector to access only the probes, and set $\rho = \tau_\beta$, $\sigma = \tau_{\beta_0}$. Equation (12) requires extra variance in every unbiased fully-PPT readout of $A_\star$. The task compares two finite temperatures; it does not claim a collective gain in local thermometric Fisher information.

For Fig. 2, we set $\beta\omega = 3$, $\beta_0\omega = 6$, and $\Omega/\omega = 1/5$. To obtain the red curve's *product-readout variance* V_{prod} , choose a Pauli axis independently on each probe and assign unbiased scores to their joint record. The scores minimize variance for this specified POVM [54]; we have not optimized over all product measurements. At each

coupling, the bounds and the feasible readout obey

$$\frac{1}{Q_2 - 1} \leq \frac{1}{Q_1^2 - 1} \leq b_{\text{mom}}(A_\star; \sigma^{\otimes 2}) \leq V_{\text{prod}}. \quad (19)$$

At the illustrative point $g/\omega = 0.25$, the spectral variance is 20.0138, the PPT moment floor is 49.5198, and the specified product readout has variance 195.8200. Hence the unknown fully-PPT minimum lies between 49.5198 and 195.8200; at least 29.5060 is added to the intrinsic variance. This receiver is one feasible upper bound, not an optimized implementation. The Pearson surplus $Q_2 - Q_1^2 \simeq 6.56 \times 10^{-4}$ is small; the stronger moment floor uses the full fixed operator. A score near 1022 has reference probability 2.26×10^{-6} , so about 80% of 10^5 -shot batches miss that outcome. These are population calculations, not a receiver validation or shot budget.

If the detector also receives each ancilla, the complete Gibbs states commute; the accessible states and target then change. This does not alter the restriction on reading the original probe score. Supplemental Sec. S6 gives the states and checks.

The examples compare complete driven-medium copies with reduced probes from separable pairs. Both use independent sources. Their accessible states do not commute, so unrestricted joint readout beats the full PPT class, including local protocols, in the measured-Pearson task.

Supplemental Sec. S7 maps the population calculations to the code and explains the finite-shot limits.

Experimental use requires calibrating the operator identity $A_\star = \sum_x a_x E_x$, not just its two endpoint means. The rare outcome makes finite-shot inference difficult. The red curve is a population upper bound, not a sampling budget. A sample variance below the green bound would require uncertainty analysis; it would not overturn the operator inequality. We report calculations without a device demonstration.

Conclusions.— For product hypotheses, the fully-PPT optimum factorizes and is attained by separate score measurements. Noncommuting states still give a strict unrestricted two-copy gain. Every unbiased fully-PPT readout of the corresponding collective score adds variance; the stronger moment bound is necessary, and the exact fixed-target minimum remains open. Dephasing can erase the gain without changing work statistics. Reduced equilibrium probes show it without a drive or entangled source. With characterized inputs and calibrated probabilities, a Pearson value above the product limit certifies an NPT readout effect.

-
- [1] K. Temme and F. Verstraete, Quantum chi-squared and goodness of fit testing, *J. Math. Phys.* **56**, 012202 (2015), arXiv:1112.6343.
- [2] Y. Tsuda and K. Matsumoto, Quantum estimation for non-differentiable models, *J. Phys. A: Math. Gen.* **38**, 1593 (2005), arXiv:quant-ph/0207150.
- [3] M. Gessner and A. Smerzi, Hierarchies of frequentist bounds for quantum metrology: From Cramér–Rao to Barankin, *Phys. Rev. Lett.* **130**, 260801 (2023), arXiv:2303.06108.
- [4] Maestro v0.1, Two-copy onset of strict superadditivity for measured relative entropy (2026), pudim AI Research Object RO-2026-0003, v1.0; program managed by D. S. P. Salazar.
- [5] T. Rippchen, S. Sreekumar, and M. Berta, Locally-measured rényi divergences, *IEEE Trans. Inf. Theory* **71**, 6105 (2025), arXiv:2405.05037.
- [6] M. Mosonyi and F. Hiai, Test-measured Rényi divergences, *IEEE Trans. Inf. Theory* **69**, 1074 (2023), arXiv:2201.05477.
- [7] Z. Huang and M. M. Wilde, Semi-definite optimization of the measured relative entropies of quantum states and channels (2024), arXiv:2406.19060.
- [8] B. M. Terhal, D. P. DiVincenzo, and D. W. Leung, Hiding bits in Bell states, *Phys. Rev. Lett.* **86**, 5807 (2001), arXiv:quant-ph/0011042.
- [9] T. Eggeling and R. F. Werner, Hiding classical data in multipartite quantum states, *Phys. Rev. Lett.* **89**, 097905 (2002), arXiv:quant-ph/0203004.
- [10] G. M. D’Ariano, P. Perinotti, and M. F. Sacchi, Informationally complete measurements on bipartite quantum systems: Comparing local with global measurements, *Phys. Rev. A* **72**, 042108 (2005), arXiv:quant-ph/0507104.
- [11] M. Berta, O. Fawzi, and M. Tomamichel, On variational expressions for quantum relative entropies, *Lett. Math. Phys.* **107**, 2239 (2017), arXiv:1512.02615.
- [12] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, *Phys. Rev. Lett.* **114**, 158101 (2015), arXiv:1502.05944.
- [13] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, Dissipation bounds all steady-state current fluctuations, *Phys. Rev. Lett.* **116**, 120601 (2016), arXiv:1512.02212.
- [14] Y. Hasegawa, Quantum thermodynamic uncertainty relation for continuous measurement, *Phys. Rev. Lett.* **125**, 050601 (2020), arXiv:1911.11982.
- [15] Y. Hasegawa, Thermodynamic uncertainty relation for general open quantum systems, *Phys. Rev. Lett.* **126**, 010602 (2021), arXiv:2003.08557.
- [16] D. S. P. Salazar, Quantum relative entropy uncertainty relation, *Phys. Rev. E* **109**, L012103 (2024), arXiv:2309.09990.
- [17] J. M. Hammersley, On estimating restricted parameters, *J. R. Stat. Soc. B* **12**, 192 (1950).
- [18] D. G. Chapman and H. Robbins, Minimum variance estimation without regularity assumptions, *Ann. Math. Stat.* **22**, 581 (1951).
- [19] A. Dechant and S. ichi Sasa, Fluctuation-response inequality out of equilibrium (2020), arXiv:1804.08250.
- [20] Y. Hasegawa and T. V. Vu, Uncertainty relations in stochastic processes: An information inequality approach, *Phys. Rev. E* **99**, 062126 (2019), arXiv:1809.03292.
- [21] M. Ozawa, Uncertainty relations for noise and disturbance in generalized quantum measurements, *Annals of Physics* **311**, 350 (2004), arXiv:quant-ph/0307057.
- [22] A. Peres, Separability criterion for density matrices, *Phys. Rev. Lett.* **77**, 1413 (1996), arXiv:quant-ph/9604005.
- [23] M. Horodecki, P. Horodecki, and R. Horodecki, Separability of mixed states: necessary and sufficient conditions, *Phys. Lett. A* **223**, 1 (1996), arXiv:quant-ph/9605038.
- [24] C. H. Bennett, D. P. DiVincenzo, C. A. Fuchs, T. Mor, E. Rains, P. W. Shor, J. A. Smolin, and W. K. Wootters, Quantum nonlocality without entanglement, *Phys. Rev. A* **59**, 1070 (1999), arXiv:quant-ph/9804053.
- [25] W. Matthews, S. Wehner, and A. Winter, Distinguishability of quantum states under restricted families of measurements with an application to quantum data hiding, *Commun. Math. Phys.* **291**, 813 (2009), arXiv:0810.2327.
- [26] E. Chitambar, D. Leung, L. Mančinska, M. Ozols, and A. Winter, Everything you always wanted to know about LOCC (but were afraid to ask), *Commun. Math. Phys.* **328**, 303 (2014), arXiv:1210.4583.
- [27] M. Piani, Relative entropy of entanglement and restricted measurements, *Phys. Rev. Lett.* **103**, 160504 (2009), arXiv:0904.2705.
- [28] F. G. S. L. Brandão, A. W. Harrow, J. R. Lee, and Y. Peres, Adversarial hypothesis testing and a quantum Stein’s lemma for restricted measurements (2020), arXiv:1308.6702.
- [29] C. W. Helstrom, *Quantum Detection and Estimation Theory* (Academic Press, New York, 1976).
- [30] A. S. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (North-Holland, Amsterdam, 1982).
- [31] S. L. Braunstein and C. M. Caves, Statistical distance and the geometry of quantum states, *Phys. Rev. Lett.* **72**, 3439 (1994).
- [32] M. G. A. Paris, Quantum estimation for quantum technology, *Int. J. Quantum Inf.* **7**, 125 (2009), arXiv:0804.2981.
- [33] D. S. P. Salazar, Benign projective landscapes for measured quantum divergences (2026), arXiv:2609.27767.
- [34] B. Regula and M. Tomamichel, Rethinking quantum smooth entropies: Tight one-shot analysis of quantum privacy amplification (2026), version 4, revised 1 September 2026, arXiv:2603.04493.
- [35] K. Temme, M. J. Kastoryano, M. B. Ruskai, M. M. Wolf, and F. Verstraete, The χ^2 -divergence and mixing times of quantum Markov processes, *J. Math. Phys.* **51**, 122201 (2010), corrected arXiv version 3 (2024), arXiv:1005.2358.
- [36] D. Petz, Quasi-entropies for finite quantum systems, *Rep. Math. Phys.* **23**, 57 (1986).
- [37] D. Petz, Monotone metrics on matrix spaces, *Linear Algebra Appl.* **244**, 81 (1996).
- [38] A. Lesniewski and M. B. Ruskai, Monotone Riemannian metrics and relative entropy on noncommutative probability spaces, *J. Math. Phys.* **40**, 5702 (1999), arXiv:math-ph/9808016.
- [39] Y. Hasegawa, Dynamics-independent bounds on state transformations and precision in open quantum systems (2025), arXiv:2508.13884.
- [40] D. S. P. Salazar, Universal thermodynamic uncertainty relation for quantum f -divergences, *Phys. Rev. E* **113**, 064150 (2026), arXiv:2511.10817.
- [41] See Supplemental Material at [URL will be inserted by publisher] for detailed proofs, model parameters, and numerical checks (2026).

- [42] H. P. Robertson, The uncertainty principle, *Phys. Rev.* **34**, 163 (1929).
- [43] D. S. P. Salazar, Matrix thermodynamic uncertainty relation for non-Abelian charge transport, *Phys. Rev. E* **114**, L022106 (2026), arXiv:2512.24956.
- [44] G. M. D’Ariano, V. Giovannetti, and P. Perinotti, Optimal estimation of quantum observables, *J. Math. Phys.* **47**, 022102 (2006), arXiv:quant-ph/0510187.
- [45] A. Miranowicz, M. Piani, P. Horodecki, and R. Horodecki, Inseparability criteria based on matrices of moments, *Phys. Rev. A* **80**, 052303 (2009), arXiv:quant-ph/0605001.
- [46] P. Talkner, E. Lutz, and P. Hänggi, Fluctuation theorems: Work is not an observable, *Phys. Rev. E* **75**, 050102 (2007), arXiv:cond-mat/0703189.
- [47] H. Spohn, Entropy production for quantum dynamical semigroups, *J. Math. Phys.* **19**, 1227 (1978).
- [48] G. T. Landi and M. Paternostro, Irreversible entropy production: From classical to quantum, *Rev. Mod. Phys.* **93**, 035008 (2021), arXiv:2009.07668.
- [49] I. A. Picatoste, A. Colla, and H.-P. Breuer, Local and global approaches to the thermodynamics of pure decoherence processes in open quantum systems, *Phys. Rev. A* **112**, 022210 (2025), arXiv:2506.11633.
- [50] J. P. Santos, L. C. Céleri, G. T. Landi, and M. Paternostro, The role of quantum coherence in non-equilibrium entropy production, *npj Quantum Inf.* **5**, 23 (2019), arXiv:1707.08946.
- [51] A. D. Varizi, A. P. Vieira, C. Cormick, R. C. Drumond, and G. T. Landi, Quantum coherence and criticality in irreversible work, *Phys. Rev. Research* **2**, 033279 (2020), arXiv:2004.00616.
- [52] H. J. D. Miller and J. Anders, Energy-temperature uncertainty relation in quantum thermodynamics, *Nat. Commun.* **9**, 2203 (2018), arXiv:1801.08057.
- [53] A. Ullah, M. T. Naseem, and Ö. E. Müstecaplıoğlu, Low-temperature quantum thermometry boosted by coherence generation, *Phys. Rev. Research* **5**, 043184 (2023), arXiv:2211.05461.
- [54] G. M. D’Ariano and P. Perinotti, Optimal data processing for quantum measurements, *Phys. Rev. Lett.* **98**, 020403 (2007), arXiv:quant-ph/0610058.

Supplemental Material for An Exact Product Limit on Collective Quantum Readout

Codex 5.6 Sol

Pudim Physics Research Ledger

Program Manager: Domingos S. P. Salazar

September 29, 2026

Where does the readout gap enter, and what does each plotted curve represent? Sections S1–S3 derive the one-copy score, the fully-PPT product limit, and the two-copy surplus. Section S4 turns these results into bounds for one prescribed observable. Sections S5 and S6 give the work-stroke and equilibrium calculations; S7 maps the figures to the data and states what finite-shot samples can show. Throughout, the hypotheses factor across the registers used to define PPT, and each comparison between readout classes holds its observable fixed.

S1 The known uncertainty identity and optimal score

S1.1 Definitions and exact single-copy solution

Let $\rho, \sigma > 0$ be density matrices on a finite-dimensional Hilbert space $\mathcal{H}$, and let $\mathcal{M}(\mathcal{H})$ denote all POVMs on that space. A POVM $M = \{E_x\} \in \mathcal{M}(\mathcal{H})$ has $E_x \succeq 0$ and $\sum_x E_x = I$. It yields $p_x = \text{Tr}[\rho E_x]$ and $q_x = \text{Tr}[\sigma E_x]$. Zero effects can be omitted; otherwise full-rankness implies $q_x > 0$. Define

$$Q(M) = \sum_x p_x^2 / q_x, \quad \mathcal{Q}(\rho, \sigma) = \max_{M \in \mathcal{M}(\mathcal{H})} Q(M). \quad (\text{S1})$$

Here $Q(M) \equiv Q(M; \rho, \sigma)$ belongs to one fixed measurement. The unrestricted optimum is $\mathcal{Q}$. For a measurement class $\mathcal{C}$, define $\mathcal{Q}^{\mathcal{C}}(\rho, \sigma) := \max_{M \in \mathcal{C}} Q(M; \rho, \sigma)$; these maxima are attained. With n identical input registers, write $\mathcal{M}_n = \mathcal{M}(\mathcal{H}^{\otimes n})$ and $Q_n(M) = Q(M; \rho^{\otimes n}, \sigma^{\otimes n})$. The class optimum is $\mathcal{Q}_n^{\mathcal{C}} := \max_{M \in \mathcal{C}_n} Q_n(M)$, with $Q_n \equiv \mathcal{Q}_n^{\mathcal{M}}$. Thus the argument of Q_n identifies a measurement, while a superscript identifies the class over which we optimize. A restriction changes that class, not the detector functional. The logarithm of $\mathcal{Q}$ is the standard measured order-two Rényi divergence, but we keep Q as the readout variable. For identical inputs, $Q_n = \mathcal{Q}(\rho^{\otimes n}, \sigma^{\otimes n})$ and $Q_0 = 1$. Temme and Verstraete [1], Lemma 5 and Eq. (44) in the arXiv PDF, establish the measured Bures χ^2 equality and the attaining score PVM. General variational expressions and projective sufficiency appear in Ref. [2]; recent smooth-entropy work also uses this quadratic quantity [3]. We give the short derivation to fix our conventions.

Let $\mathcal{J}_\sigma(A) = (\sigma A + A \sigma)/2$ on the Hilbert–Schmidt space. It is self-adjoint and positive definite. Its inverse maps positive matrices to positive matrices because

$$\mathcal{J}_\sigma^{-1}(B) = 2 \int_0^\infty e^{-t\sigma} B e^{-t\sigma} dt. \quad (\text{S2})$$

In particular, $L = \mathcal{J}_\sigma^{-1}(\rho) > 0$ is well defined. For any real outcome scores a_x ,

$$\sum_x \frac{p_x^2}{q_x} = \max_{\{a_x\}} \sum_x (2p_x a_x - q_x a_x^2). \quad (\text{S3})$$

Set $A = \sum_x a_x E_x$. The Schwarz inequality for the unital POVM map gives $A^2 \preceq \sum_x a_x^2 E_x$. The same inequality bounds $Q(M)$ above by the maximum of

$$F(A) = 2 \text{Tr}[\rho A] - \text{Tr}[\sigma A^2], \quad A = A^\dagger. \quad (\text{S4})$$

Since $\mathcal{J}_\sigma(L) = \rho$,

$$F(L) - F(A) = \text{Tr}[\sigma(A - L)^2], \quad \mathcal{Q} = F(L) = \text{Tr}[\rho L] = \text{Tr}[\sigma L^2]. \quad (\text{S5})$$

To see that the upper bound is attainable, use the spectral decomposition $L = \sum_x \ell_x P_x$. The projective measurement $\{P_x\}$ reports the eigenvalue ℓ_x ; each P_x projects onto its eigenspace. If $LP_x = \ell_x P_x$, then $\text{Tr}[\rho P_x] = \text{Tr}[\mathcal{J}_\sigma(L) P_x] = \ell_x \text{Tr}[\sigma P_x]$. The optimal classical score is therefore $p_x/q_x = \ell_x$, and this PVM attains Eq. (S5). Repeated score eigenvalues cause no problem. In an eigenbasis of σ ,

$$L_{ij} = \frac{2\rho_{ij}}{s_i + s_j}, \quad \mathcal{Q} = \sum_{ij} \frac{2|\rho_{ij}|^2}{s_i + s_j}. \quad (\text{S6})$$

No new theorem is assigned to these identities.

S1.2 Sharp Hammersley–Chapman–Robbins form

For a Hermitian observable A , put $A_c = A - \text{Tr}[\sigma A] I$. In the real inner product $\langle B, C \rangle_\sigma = \text{Tr}[B \mathcal{J}_\sigma(C)]$ on Hermitian matrices,

$$\text{Tr}[(\rho - \sigma)A] = \langle L - I, A_c \rangle_\sigma. \quad (\text{S7})$$

Cauchy–Schwarz and $\text{Tr}[\sigma L] = 1$ give

$$\frac{[\text{Tr}[(\rho - \sigma)A]]^2}{\text{Var}_\sigma A} \leq \text{Tr}[\sigma(L - I)^2] = \mathcal{Q} - 1. \quad (\text{S8})$$

Equality holds for $A = L - I$, up to nonzero scaling and a scalar shift. For $\rho = \sigma$, the optimal ratio is zero. Optimizing the numerical score of a fixed POVM gives its classical Pearson divergence $Q(M) - 1$ by the same argument. Thus the receiver objective is the optimal finite-displacement signal-to-reference-variance ratio, plus one. It is not the Bayesian error probability of discriminating the hypotheses.

Classical finite-difference estimation bounds go back to Hammersley and Chapman–Robbins [4, 5]. Tsuda and Matsumoto’s quantum Hammersley–Chapman–Robbins theorem [6], Theorem 1, defines the finite-difference SLD through $\mathcal{J}_{\rho_\theta}(L_{\theta,\delta}^S) = (\rho_{\theta+\delta} - \rho_\theta)/\delta$. Take $\rho_\theta = \sigma$, $\rho_{\theta+\delta} = \rho$, and $\delta = 1$. Then $L^S = L - I$, and their SLD denominator is exactly $\mathcal{Q} - 1$. The spectral measurement gives the corresponding saturation. The same denominator is the minimal (Bures/SLD) quantum χ^2 divergence in Ref. [7]. We use the corrected arXiv version of that work and do not rely on its removed earlier Proposition 6.

S2 The fully-PPT bound and its attaining measurement

Consider N registers with Hilbert space $\bigotimes_{j=1}^N \mathcal{H}_j$. Write $\mathcal{M}_N$ for all POVMs on this space and $[N] = \{1, \dots, N\}$. Fix one orthonormal basis in each register. For $S \subseteq [N]$, $T_S(E) = E^{Ts}$ transposes the matrix indices of the registers in S and leaves the others unchanged. In the product basis grouped as $S : \bar{S}$, it acts linearly as

$$(|a\rangle\langle c| \otimes |b\rangle\langle d|)^{Ts} = |c\rangle\langle a| \otimes |b\rangle\langle d|, \quad \langle ab|E^{Ts}|cd\rangle = \langle cb|E|ad\rangle.$$

Thus T_S is neither an exponent nor a physical transpose gate. The condition $E^{Ts} \succeq 0$ means that the Hermitian matrix E^{Ts} has only nonnegative eigenvalues, not that its entries are nonnegative. Zero eigenvalues are allowed. For $M = \{E_x\} \in \mathcal{M}_N$, define

$$M \in \mathcal{P}_N \iff E_x^{Ts} \succeq 0 \text{ for every } x \text{ and every } S \subseteq [N]. \quad (\text{S9})$$

Ordinary positivity $E_x \succeq 0$ is already part of the POVM definition. The empty subset adds nothing, and $T_{[N]}$ is the full transpose, which preserves positivity. For $N = 1$ the restriction is vacuous: $\mathcal{P}_1 = \mathcal{M}_1$. For $N = 2$ either nontrivial subset gives the usual bipartite PPT condition. For more registers one must test every cut, not only a chosen bipartition. Changing the fixed local bases conjugates a partial transpose by a unitary, so positivity does not depend on that choice.

Fully separable effects and LOCC measurements meet this condition [8–10], but fully PPT is not a synonym for separable or classical. For identical inputs the numerical optimum is $Q_n^{\mathcal{P}} := \max_{M \in \mathcal{P}_n} Q_n(M)$. Thus $Q_n^{\mathcal{P}} = \mathcal{Q}^{\mathcal{P}_n}(\rho^{\otimes n}, \sigma^{\otimes n})$. The Letter’s $\mathcal{Q}^{\mathcal{P}_N}$ uses the specified register partition also for heterogeneous pairs. The hypotheses used here factor over those same registers. Internal subsystems of a register are not separate PPT factors.

For local full-rank pairs, set

$$L_j = \mathcal{J}_{\sigma_j}^{-1}(\rho_j), \quad Y_j = L_j \sigma_j L_j, \quad \sigma^{(N)} = \bigotimes_j \sigma_j, \quad \rho^{(N)} = \bigotimes_j \rho_j. \quad (\text{S10})$$

For each subset S , transpose the local factors in both L_j and σ_j :

$$\tilde{L}_S = \bigotimes_j L_j^{T^{1_{j \in S}}}, \quad \tilde{\sigma}_S = \bigotimes_j \sigma_j^{T^{1_{j \in S}}}. \quad (\text{S11})$$

Here T^0 means no transpose and T^1 means transpose. For any real u we have the exact identity

$$\bigotimes_j Y_j - 2u\rho^{(N)} + u^2\sigma^{(N)} = 2^{-N} \sum_S T_S \left[(\tilde{L}_S - uI) \tilde{\sigma}_S (\tilde{L}_S - uI) \right]. \quad (\text{S12})$$

To verify it, expand each bracket. The quadratic and constant terms become respectively $\bigotimes_j Y_j$ and $u^2\sigma^{(N)}$, independent of S . In the cross terms, partial transposition reverses the local order on S . Averaging all subsets produces every ordering in the expansion of $\bigotimes_j (\sigma_j L_j + L_j \sigma_j)/2 = \rho^{(N)}$.

Every matrix inside the square brackets is positive. For a fully-PPT effect E and the numbers

$$p = \text{Tr} \left[E \rho^{(N)} \right], \quad q = \text{Tr} \left[E \sigma^{(N)} \right], \quad y = \text{Tr} \left[E \left(\bigotimes_j Y_j \right) \right], \quad (\text{S13})$$

Eq. (S12) implies $y - 2up + u^2q \geq 0$ for all real u . Full rank gives $q > 0$ for a nonzero effect. Minimizing over u proves $p^2/q \leq y$. Summation over a POVM and $\sum_x E_x = I$ now give

$$Q(M) \leq \text{Tr} \left[\bigotimes_j Y_j \right] = \prod_j \text{Tr} \left[\sigma_j L_j^2 \right] = \prod_j \mathcal{Q}(\rho_j, \sigma_j), \quad M \in \mathcal{P}_N. \quad (\text{S14})$$

Conversely, the product of the local score PVMs is separable and reaches the product value. This proves main-text Theorem 1 for heterogeneous states and dimensions. Section S4 converts this value into the attained uncertainty coefficient and proves the fixed-observable noise bound.

S2.1 Measurement scope and detector interpretation

Theorem 1 does not identify PPT with separability. Its upper bound applies to the larger fully-PPT class, while a product measurement establishes achievability. The bound therefore applies beyond independent readout. The decomposition is finite and exact. It needs no optimization-duality assumption and no closure property under operator powers.

Suppose a calibrated detector gives $Q(M) > \prod_j Q_j$. At least one outcome operator must then have a negative partial transpose. The converse need not hold: an NPT measurement can fail to improve on

product readout. Failure to exceed the bound proves neither that the detector is PPT nor that the states commute. This test assumes known input states.

The state assumption and partition differ from the locally measured Rényi examples in Ref. [11]. Their isotropic and Werner constructions can be correlated or entangled inside each $A : B$ pair and their many-copy measurement cut can be $A^n : B^n$. Here the hypotheses are products *across the registers on which PPT is imposed*. The distinction prevents importing a data-hiding formula for a different partition.

S3 Ordered commutator surplus, stability, and finite-displacement onset

S3.1 Exact two-source residual

For a local pair define the Hermitian matrix $K = [\sigma, L]/(2i)$. Then $\sigma L = \rho + iK$ and $L\sigma = \rho - iK$. For two local pairs this gives

$$\mathcal{J}_{\sigma_A \otimes \sigma_B}(L_A \otimes L_B) = \rho_A \otimes \rho_B - K_A \otimes K_B. \quad (\text{S15})$$

In the joint quadratic variational problem the trial $A_0 = L_A \otimes L_B$ has value $Q_A Q_B$, because

$$2 \text{Tr}[(\rho_A \otimes \rho_B)A_0] - \text{Tr}[(\sigma_A \otimes \sigma_B)A_0^2] = 2Q_A Q_B - Q_A Q_B. \quad (\text{S16})$$

For the exact joint score L_{AB} , Eq. (S15) says $L_{AB} - A_0 = \mathcal{J}_{\sigma_A \otimes \sigma_B}^{-1}(K_A \otimes K_B)$. Completing the square with Eq. (S5) yields

$$\Delta_{AB} := Q_{AB} - Q_A Q_B = \langle K_A \otimes K_B, \mathcal{J}_{\sigma_A \otimes \sigma_B}^{-1}(K_A \otimes K_B) \rangle_{\text{HS}}. \quad (\text{S17})$$

The inverse is strictly positive as a quadratic form, so equality holds precisely when $K_A \otimes K_B = 0$. In a reference eigenbasis, $K_{ij} = (s_i - s_j)\rho_{ij}/[i(s_i + s_j)]$, which vanishes precisely when $[\rho, \sigma] = 0$. This proves main-text Proposition 1 without a nondegeneracy assumption.

S3.2 Positive ordered-moment expansion

For each local pair put

$$w_{ij} = \frac{2|\rho_{ij}|^2}{s_i + s_j}, \quad t_{ij} = \frac{s_i - s_j}{s_i + s_j}, \quad A_m = \sum_{ij} w_{ij} t_{ij}^{2m} \quad (m \geq 1). \quad (\text{S18})$$

The denominator of a tensor term factorizes as

$$s_i s_k + s_j s_l = \frac{1}{2}(s_i + s_j)(s_k + s_l)(1 + t_{ij} t_{kl}). \quad (\text{S19})$$

Swapping i, j leaves w_{ij} fixed and reverses t_{ij} . Therefore symmetrizing the two terms replaces $(1 + t_{ij} t_{kl})^{-1}$ by $(1 - t_{ij}^2 t_{kl}^2)^{-1}$. The uniformly convergent geometric series, valid since all $s_i > 0$, gives

$$\Delta_{AB} = \sum_{m=1}^{\infty} A_m^A A_m^B. \quad (\text{S20})$$

The first coefficient and a useful finite estimate are

$$A_1 = 2 \sum_{ij} \frac{|[\rho, \sigma]_{ij}|^2}{(s_i + s_j)^3}, \quad (\text{S21})$$

$$A_1^A A_1^B \leq \Delta_{AB} \leq \frac{A_1^A A_1^B}{1 - \eta_A^2 \eta_B^2}, \quad \eta = \frac{s_{\max} - s_{\min}}{s_{\max} + s_{\min}} < 1. \quad (\text{S22})$$

These bounds do not imply a state-independent positive minimum advantage.

Let $X_{ij} = |i\rangle\langle j| + |j\rangle\langle i|$ and $Y_{ij} = -i|i\rangle\langle j| + i|j\rangle\langle i|$. Their ordered covariance in σ is

$$C_{\sigma}^{ij} = \begin{pmatrix} s_i + s_j & i(s_i - s_j) \\ -i(s_i - s_j) & s_i + s_j \end{pmatrix}. \quad (\text{S23})$$

Thus t_{ij} is the imaginary covariance divided by the diagonal variance. Equations (S20)–(S22) express the surplus as positive contributions from this ordering information. This connects the calculation to ordered covariance [12, 13] without asserting that every quantum uncertainty bound is saturated.

S4 Readout variance and added fluctuations of a prescribed observable

We now fix the two preparations and the observable while changing the readout. The observable's intrinsic variance stays the same. The variance of the reported outcomes can change because the measurement may add noise.

S4.1 First and second moments of an unbiased readout

Let $R, S > 0$ be distinct density matrices in finite dimension, and let $M = \{E_x\}$ be a POVM. A real output assignment a_x is an unbiased readout of $A = A^\dagger$ when

$$A = \sum_x a_x E_x, \quad B = \sum_x a_x^2 E_x. \quad (\text{S24})$$

The first equality enforces the correct expectation for every input state. It is stronger than calibrating only two expectation values. No assertion is made that the complete output distribution equals the spectral distribution of A .

Introduce the isometry $W\psi = \bigoplus_x \sqrt{E_x} \psi$ and the block-diagonal operator $D = \bigoplus_x a_x I$. Then $W^\dagger W = I$, $A = W^\dagger D W$ and $B = W^\dagger D^2 W$. Therefore

$$N := B - A^2 = W^\dagger D(I - W W^\dagger) D W \succeq 0. \quad (\text{S25})$$

For this pair define $q_x = \text{Tr}[S E_x]$, $j(A) = \text{Tr}[(R - S)A]$, and $v(A) = \text{Var}_S A$. Thus $j(A) = \langle A \rangle_R - \langle A \rangle_S$ is a mean response, not necessarily a transport current. We obtain

$$V_S(M, a) = \sum_x q_x a_x^2 - (\text{Tr}[S A])^2 = v(A) + \nu_S(M, a; A), \quad \nu_S = \text{Tr}[S N] \geq 0. \quad (\text{S26})$$

The intrinsic noise operator in Eq. (S25) is standard measurement theory [14]. Under the imposed first-moment equality its expectation is added readout variance. It is not a disturbance of a second observable or an energetic cost. Spectral readout of A gives $N = 0$. Since $S > 0$, $\nu_S = 0$ implies $N = 0$.

S4.2 A sharp class-wide coefficient, and its fixed-observable consequence

For $p_x = \text{Tr}[R E_x]$ and $z_x = p_x/q_x - 1$, normalization gives $\sum_x q_x z_x = 0$. Hence

$$j(A) = \sum_x q_x z_x (a_x - \text{Tr}[S A]), \quad j(A)^2 \leq [Q(M; R, S) - 1] V_S(M, a). \quad (\text{S27})$$

For $Q(M) > 1$, equality holds exactly when $a_x - \text{Tr}[S A]$ is a real constant times z_x on the nonzero effects; the constant is nonzero when $j(A) \neq 0$. When $Q(M) = 1$, every recorded contrast is zero.

Let $\mathcal{C}$ be a measurement class with attained optimum $Q^c(R, S) > 1$. Equation (S27) implies, for every unbiased readout of any prescribed A ,

$$V_S(M, a) \geq \frac{j(A)^2}{Q^c - 1}, \quad \nu_S(M, a; A) \geq \left[\frac{j(A)^2}{Q^c - 1} - \text{Var}_S A \right]_+, \quad M \in \mathcal{C}. \quad (\text{S28})$$

There is no requirement that A be a current, energy, or work observable. Ratios involving $j(A)$ assume $j(A) \neq 0$; Eq. (S27) covers zero contrast without division. If $R = S$, every contrast vanishes and the normalized optimal observable below is not defined.

To locate the equality case, fix two numerical means μ and $\mu + \delta$, with $\delta \neq 0$. Leave the full operator A free. For an attaining measurement $M_{\mathcal{C}}$ set

$$a_x = \mu + \frac{\delta}{Q^c - 1} \left(\frac{p_x}{q_x} - 1 \right). \quad (\text{S29})$$

Its means under S, R are exactly $\mu, \mu + \delta$, and its variance is $\delta^2/(Q^c - 1)$. Thus

$$\min_{\substack{M \in \mathcal{C}, \\ a \in \mathbb{R}^{|M|} \\ \sum_x q_x a_x = \mu, \sum_x p_x a_x = \mu + \delta}} V_S(M, a) = \frac{\delta^2}{Q^c - 1}. \quad (\text{S30})$$

This optimum holds when only the two means are fixed. Fixing the whole operator A can prevent equality. Equation (S28) therefore remains a lower bound on its added noise, rather than a formula for the restricted minimum of every A .

For $R = \bigotimes_j \rho_j$, $S = \bigotimes_j \sigma_j$, and the fully-PPT class across those same registers, Section S2 gives $Q^{\mathcal{P}} = \prod_j Q_j$, with $Q_j = \mathcal{Q}(\rho_j, \sigma_j)$. Its product score measurement is an attaining $M_{\mathcal{C}}$. The output in Eq. (S29) is a product likelihood ratio followed by an affine transformation. It is generally not a sum of local currents. For identical pairs the coefficient is $Q_1^n - 1$, not $n(Q_1 - 1)$. The nonlinearity already available through classical output processing is not counted as a quantum surplus.

S4.3 The same optimal observable needs extra variance under PPT readout

For $R = \rho_n$, $S = \sigma_n$, let $L_n = \mathcal{J}_S^{-1}(R)$ and define

$$A_{\star} = \frac{L_n - I}{Q_n - 1}. \quad (\text{S31})$$

From $\text{Tr}[SL_n] = 1$ and $\text{Tr}[RL_n] = \text{Tr}[SL_n^2] = Q_n$,

$$\text{Tr}[SA_{\star}] = 0, \quad \text{Tr}[RA_{\star}] = 1, \quad \text{Var}_S A_{\star} = \frac{1}{Q_n - 1}. \quad (\text{S32})$$

Unrestricted spectral readout has this variance and no added noise. Every unbiased fully-PPT readout of the *same* $A_{\star}$ satisfies

$$\nu_S(M, a; A_{\star}) \geq \frac{1}{Q_1^n - 1} - \frac{1}{Q_n - 1} = \frac{Q_n - Q_1^n}{(Q_1^n - 1)(Q_n - 1)}. \quad (\text{S33})$$

Equivalently,

$$\frac{\nu_S(M, a; A_{\star})}{\text{Var}_S A_{\star}} \geq \frac{Q_n - Q_1^n}{Q_1^n - 1}. \quad (\text{S34})$$

For noncommuting states and $n \geq 2$ the right-hand sides are positive by Section S3. Changing the readout cannot change the first moments or intrinsic variance in Eq. (S32); it can only change the

output second moment subject to Eq. (S33). For the unnormalized score L_n , multiplication by $(Q_n - 1)^2$ yields

$$\nu_S(M, a; L_n) \geq \frac{(Q_n - 1)(Q_n - Q_1^n)}{Q_1^n - 1}. \quad (\text{S35})$$

A product informationally complete POVM spans the Hermitian operator space, so real scores can reconstruct $A_\star$ exactly in first moment. An unbiased restricted readout therefore exists. Its variance still obeys the penalty. The script `verify_readout_tur.py` reconstructs the scores and checks the inequality; it does not optimize over all PPT POVMs.

For contrast, $A = \sum_j A_j^{(j)}$ has a product spectral implementation. Its output second moment equals A^2 , giving $\nu_S = 0$. This is consistent with Eq. (S28), whose right-hand side must vanish for this target. The theorem is not a claim of strict advantage for every physical observable.

S4.4 A partial-transpose moment refinement

Let $S = \bigotimes_{j=1}^n \sigma_j > 0$ be the reference state and fix a Hermitian A on the full register space. An unbiased readout obeys $A = \sum_x a_x E_x$ with real scores. Put $B = \sum_x a_x^2 E_x$ and $\mu = \text{Tr}[SA]$. For each subset $J \subseteq [n]$ define

$$v_J(A; S) = \text{Var}_{S^{T_J}}(A^{T_J}), \quad b_{\text{mom}}(A; S) = \max_{J \subseteq [n]} v_J(A; S). \quad (\text{S36})$$

The empty subset is included, so $b_{\text{mom}} \geq \text{Var}_S A$. A transposition is taken in fixed local bases; simultaneous local changes of basis do not change this number.

Proposition S1 (Fixed-observable PPT moment bound) *Every unbiased fully-PPT readout of A satisfies*

$$V_S(M, a) \geq b_{\text{mom}}(A; S), \quad \nu_S(M, a; A) \geq b_{\text{mom}}(A; S) - \text{Var}_S A. \quad (\text{S37})$$

This bound does not require a signal state. For a product signal $R = \bigotimes_j \rho_j$, it dominates the two-mean readout bound:

$$b_{\text{mom}}(A; S) \geq \bar{v}(A; S) \geq \frac{[\text{Tr}[(R - S)A]]^2}{\prod_j \mathcal{Q}(\rho_j, \sigma_j) - 1}, \quad \bar{v} = 2^{-n} \sum_{J \subseteq [n]} v_J. \quad (\text{S38})$$

The last ratio assumes $R \neq S$.

Proof. For $M \in \mathcal{P}_n$, each family $\{E_x^{T_J}\}$ is a POVM: its effects are positive and sum to I . Applying the ordinary POVM Schwarz inequality to this family gives

$$B^{T_J} \succeq (A^{T_J})^2. \quad (\text{S39})$$

Since $S^{T_J} > 0$, trace pairing with it and the self-adjointness and involution of partial transposition give

$$\text{Tr}[SB] = \text{Tr}[S^{T_J} B^{T_J}] \geq \text{Tr}[S^{T_J} (A^{T_J})^2].$$

Also $\text{Tr}[S^{T_J} A^{T_J}] = \mu$. Subtracting μ^2 and maximizing over J proves Eq. (S37). The same first step applies to any reference positive under the specified partial transposes, but the following factorization uses a product reference.

Set $A_c = A - \mu I$ and $\mathcal{K} = \bigotimes_j \mathcal{J}_{\sigma_j}$. Averaging all subsets yields the exact identity

$$\bar{v}(A; S) = \text{Tr}[A_c \mathcal{K}(A_c)]. \quad (\text{S40})$$

For example, diagonalize each σ_j . With multi-indices $\mathbf{a}, \mathbf{b}$, both sides equal

$$\sum_{\mathbf{a}, \mathbf{b}} |(A_c)_{\mathbf{a}\mathbf{b}}|^2 \prod_j \frac{s_{a_j}^{(j)} + s_{b_j}^{(j)}}{2}.$$

Each local index transposition exchanges which eigenvalue weights that entry; averaging generates the product of the two local choices.

Let $L_j = \mathcal{J}_{\sigma_j}^{-1}(\rho_j)$ and $L_\Pi = \bigotimes_j L_j$. Then $\mathcal{K}(L_\Pi) = R$ and $\mathcal{K}(I) = S$. In the positive inner product $\langle C, D \rangle_{\mathcal{K}} = \text{Tr}[C\mathcal{K}(D)]$ on Hermitian matrices,

$$\begin{aligned} \text{Tr}[(R - S)A] &= \langle L_\Pi - I, A_c \rangle_{\mathcal{K}}, \\ \|L_\Pi - I\|_{\mathcal{K}}^2 &= \prod_j \mathcal{Q}(\rho_j, \sigma_j) - 1. \end{aligned} \tag{S41}$$

Cauchy–Schwarz proves the second inequality in Eq. (S38); the maximum exceeds the average. This proof uses standard POVM moment positivity [14]; moment and partial-transpose methods are established also in the different context of state-separability tests [15]. No historical novelty is assigned to those ingredients. $\square$

Where the bounds agree. For $A = L_\Pi$, its product spectral readout is fully PPT and noiseless. Every v_J then equals $\prod_j Q_j - 1$, so both inequalities in Eq. (S38) are equalities. An affine rescaling gives any two prescribed distinct means. This gives another route to the sharp class-wide coefficient, while the bound for a general prescribed operator can be much larger. It does not solve the full optimization over all unbiased PPT implementations.

Checks, not a numerical proof of optimality. The suite `verify_moment_bounds.py` tests the subset-average identity on random complex states and observables, all-cut output inequalities for product readout, and equality for product optimal scores. It also tests cases beyond this shortened supplement. The recorded suite passes 1067 checks, with largest relative identity residual below 2×10^{-15} . These floating-point checks neither establish tightness of a fixed-observable bound nor replace the proof above.

S5 Work-stroke dephasing: invariant work statistics, decaying read-out surplus

S5.1 Protocol and unchanged work distribution

Prepare the medium in $\gamma_\beta = e^{-\beta H}/Z_\beta$. A cyclic pulse U leaves it in $\rho_0 = U\gamma_\beta U^\dagger$, with H restored. During the subsequent delay it obeys $\dot{\rho}_t = -i[H, \rho_t] + \Gamma(\mathcal{D}_H \rho_t - \rho_t)$. Here $\mathcal{D}_H(X) = \sum_E \Pi_E X \Pi_E$ removes coherence between distinct energy blocks. This map commutes with Hamiltonian evolution, giving the exact state

$$\rho_t = e^{-itH} \{ \mathcal{D}_H \rho_0 + c_t(\rho_0 - \mathcal{D}_H \rho_0) \} e^{itH}, \quad c_t = e^{-\Gamma t}. \tag{S42}$$

Energy populations stay constant, so the system energy and mean cyclic-pulse work do too. This system-level statement gives no value for reservoir heat or the cost of implementing the channel. The horizontal axis of main Fig. 1 is the dephasing entropy $\Sigma_\phi(t) = S(\rho_t) - S(\rho_0)$. It is nonnegative for this unital channel and differs from the collision surplus.

The standard two-projective-energy-measurement work protocol [16] has initial energy E and final energy E' . Even with degenerate energy blocks, its joint probability after the delay is

$$\Pr_t(E, E') = \frac{e^{-\beta E}}{Z_\beta} \text{Tr}[\Pi_{E'} U \Pi_E U^\dagger], \quad W = E' - E. \quad (\text{S43})$$

Every transition probability in Eq. (S43), and therefore the full work distribution, is independent of t . The timing matters: noise during the pulse may change those transitions. Work and coherent score readout use separate preparations because a final projective energy measurement would remove the coherence needed for the score task on that copy.

S5.2 Exact decay and scope of the variance bound

Set $\sigma = \gamma_\beta$ and $Q_j(t) = \mathcal{Q}(\rho_t^{\otimes j}, \sigma^{\otimes j})$. The reference state is a function of H , so terms within an energy block are orthogonal in the Jordan metric to terms between blocks. Define $Q_E = \mathcal{Q}(\mathcal{D}_H \rho_0, \sigma)$. Equation (S42) then gives

$$Q_1(t) = Q_E + c_t^2[Q_1(0) - Q_E]. \quad (\text{S44})$$

Let $L_t = \mathcal{J}_\sigma^{-1}(\rho_t)$ and $K_t = [\sigma, L_t]/(2i)$. The energy-block part of L_t commutes with σ . Jordan-map covariance under e^{-itH} therefore gives $K_t = c_t e^{-itH} K_0 e^{itH}$. The ordered-moment formula of Sec. S3 is quadratic in a pair of these commutator factors. It follows that

$$\Delta_2(t) = Q_2(t) - Q_1(t)^2 = c_t^4 \Delta_2(0). \quad (\text{S45})$$

The fourth-power law holds for this uniform-rate energy-dephasing channel, including degenerate H . It describes the difference between optimized unrestricted and fully-PPT collision values. Entropy production and the fixed-observable variance bound follow other time dependences; the latter also has changing denominators. At each delay in main Fig. 1, the readout classes are compared on the same $A_\star(t)$. That target changes along the trajectory.

S5.3 Interacting four-spin calculation

The open chain and square pulse are

$$H = -J \sum_{j=1}^3 Z_j Z_{j+1} - h \sum_{j=1}^4 X_j - b \sum_{j=1}^4 Z_j, \quad U = e^{-i\tau(H + \lambda \sum_{j=1}^4 Z_j)}. \quad (\text{S46})$$

We take $h/J = 0.8$, $b/J = 0.17$, $\lambda/J = 0.6$, $\beta J = 1$, $J\tau = 1.2$, and $\Gamma/J = 0.25$. At the start of the delay, $Q_1 = 4.9766570144$ and $Q_2 = 29.4165840530$. Their two-copy surplus is $\Delta_2 = 4.6494690142$. The unit-response target is $A_\star = (L_2 - I)/(Q_2 - 1)$. Its intrinsic variance is $1/(Q_2 - 1) = 0.0351907\dots$. The necessary fully-PPT moment bound is $0.1548592\dots$, so readout adds at least $0.1196685\dots$ in variance. These population calculations give bounds, not observed detector fluctuations or a fully-PPT implementation of this target.

The plotted points follow the exact state in Eq. (S42). The script `code/simulate_thermodynamics.py` also integrates the master equation and records every energy-transition probability before display binning. The companion `code/verify_thermodynamics.py` checks the decay and two-copy score calculation, along with energy and work invariance and the cyclic Jarzynski identity. These checks support the finite-system calculation. The medium has 16 energy levels; its optimal unrestricted two-copy score acts on 256 dimensions. We have not modeled source preparation, storage, calibration, detector cost, or a thermodynamic limit.

S6 Reduced equilibrium probes without an entangled source

The second example uses reduced equilibrium probes. Their temperature-dependent eigenvectors are already known in finite-coupling thermometry [17]. We take the one-ancilla Hamiltonian and probe marginal of Ullah, Naseem, and Müstecaplıoğlu [18]. The question here is how accurately a fixed collective observable can be read when two temperatures are compared and the detector is fully PPT.

S6.1 Exact states, separability, and noncommutation

Let $\omega, \Omega > 0$ and $g \geq 0$. Identity factors are implicit in

$$H_{SE} = \omega Z_S + \Omega Z_E + g X_S Z_E, \quad h = \sqrt{\omega^2 + g^2}. \quad (\text{S47})$$

This is Ref. [18], Eq. (1), with one ancilla, $\omega_p = 2\omega$, and $\omega_1 = 2\Omega$. Conditional on $Z_E = e \in \{-1, 1\}$, the probe Hamiltonian is $H_e = \omega Z_S + eg X_S$, with eigenvalues $\pm h$. Thus

$$\begin{aligned} G_\beta &= \frac{e^{-\beta H_{SE}}}{Z_{SE}(\beta)} = \sum_{e=\pm 1} w_e(\beta) \gamma_e(\beta) \otimes |e\rangle\langle e|, \\ w_e(\beta) &= \frac{e^{-\beta e \Omega}}{2 \cosh(\beta \Omega)}, \quad \gamma_e(\beta) = \frac{I - \tanh(\beta h)(\omega Z_S + eg X_S)/h}{2}. \end{aligned} \quad (\text{S48})$$

This is explicitly separable across $S : E$, and classical on E . Tracing out E gives

$$\tau_\beta = \frac{1}{2}[I + r_x(\beta)X_S + r_z(\beta)Z_S], \quad r_x = \frac{g}{h} \tanh(\beta h) \tanh(\beta \Omega), \quad r_z = -\frac{\omega}{h} \tanh(\beta h). \quad (\text{S49})$$

After converting energy conventions, this is the one-ancilla marginal of Ref. [18], Eqs. (13), (17)–(19). Its coherence was already described there. The global state and probe marginal are full rank at finite temperature.

For two different inverse temperatures $\beta, \beta_0 > 0$, abbreviate $t = \tanh(\beta h)$ and $t_0 = \tanh(\beta_0 h)$. Likewise, put $u = \tanh(\beta \Omega)$ and $u_0 = \tanh(\beta_0 \Omega)$. The Bloch vectors satisfy

$$(\mathbf{r}_\beta \times \mathbf{r}_{\beta_0})_y = \frac{g\omega}{h^2} t t_0 (u - u_0), \quad [\tau_\beta, \tau_{\beta_0}] = \frac{i}{2} (\mathbf{r}_\beta \times \mathbf{r}_{\beta_0})_y Y_S. \quad (\text{S50})$$

Thus the marginals fail to commute whenever $g\omega\Omega \neq 0$ and the positive finite temperatures differ. At $g = 0$, the bare-probe Gibbs states commute and the collective surplus vanishes. The effect already appears at finite coupling; no strong-coupling limit is required.

S6.2 The same observable under free and PPT readout

Prepare the probe–ancilla pairs independently and give the receiver only the probes. Set $\rho = \tau_\beta$ and $\sigma = \tau_{\beta_0}$. On two probes the hypotheses are $R = \rho^{\otimes 2}$ and $S = \sigma^{\otimes 2}$; no ancilla is shared across replicas. The product theorem and ordered-surplus identity apply to these accessible states. Choose $A_\star = (L_2 - I)/(Q_2 - 1)$. Its means are $\text{Tr}[SA_\star] = 0$ and $\text{Tr}[RA_\star] = 1$, while its intrinsic variance is $v_0 = (Q_2 - 1)^{-1}$. Spectral readout adds no variance. Every unbiased PPT readout obeys

$$V_S(M, a) \geq b_{\text{mom}}(A_\star; S) \geq \frac{1}{Q_1^2 - 1} > v_0. \quad (\text{S51})$$

At any one coupling, both readout classes see the same states and the same $A_\star$. Changing the plotted coupling changes all three. The final strict inequality requires distinct, noncommuting inputs. Neither the moment bound nor the class-wide two-mean bound determines the exact restricted minimum for this operator.

At $\beta\omega = 3$, $\beta_0\omega = 6$, $\Omega/\omega = 0.2$, and $g/\omega = 0.25$, the values are

Quantity	Value
Q_1	1.02435834934524
Q_2	1.04996555099548
v_0	20.0137891022
Two-mean readout variance lower bound	20.2798506334
PPT moment variance lower bound	49.5198347913
Added-variance lower bound from moments	29.5060456891
Moment lower bound / intrinsic variance minus one	1.47428583055
A feasible product-readout variance	195.819952797
$\lambda_{\min}(\rho), \lambda_{\min}(\sigma)$	0.0125893, 0.00451010

The moment bound is well above the intrinsic variance; the class-wide two-mean bound is much closer to it. The table uses the same intrinsic-variance denominator for its relative entry. Both lower bounds are necessary, and neither is claimed to be attained. The defining parameters and full-precision matrices determine the target; the rounded values only describe it.

S6.3 An explicit finite product implementation

A product informationally complete measurement provides a constructive upper bound on the unknown restricted minimum. Diagonalize the local reference and rotate the Pauli axes into that frame. On each probe use the six effects

$$F_{k,s} = \frac{I + s\tilde{\sigma}_k}{6}, \quad k = x, y, z, \quad s = \pm 1. \quad (\text{S52})$$

Operationally, this means choosing each of three Pauli axes with probability 1/3 and recording its binary result. Use $E_x = F_{k,s} \otimes F_{l,t}$ on the two probes. These 36 effects are product operators and span the full Hermitian space. They can read any $A_\star$ without first-moment bias, though not with its spectral distribution.

We choose the scores by the established minimum-variance data-processing prescription [19]. Use a real orthonormal basis of Hermitian matrices. Let the columns of T be the vectorized effects, and let a_{tar} represent $A_\star$ in the same basis. Put $W = \text{diag}(q_x)$, with $q_x = \text{Tr}[SE_x] > 0$. A minimum-second-moment solution is

$$a = W^{-1}T^\top(TW^{-1}T^\top)^{-1}a_{\text{tar}}. \quad (\text{S53})$$

The implementation uses a stable weighted least-squares solve rather than explicitly forming this inverse. It verifies $\sum_x a_x E_x = A_\star$ on the full four-dimensional Hilbert space, not only on the two calibrated states. The resulting variance is 195.81995..., so

$$49.51983\dots \leq \inf_{\substack{M \in \mathcal{P}_2, \, a \in \mathbb{R}^{|M|} \\ \sum_x a_x E_x = A_\star}} V_S(M, a) \leq 195.81995\dots \quad (\text{S54})$$

The lower number comes from the necessary moment inequality. The upper number comes from the specified product POVM with its optimal scores. Numerical residuals are recorded separately; they are not formal interval certificates. We have not optimized over all PPT POVMs. At the marked point, product scores range from about -24.89 to 1021.61 . Source calibration, detector-error analysis, and an experimental sampling budget would have to account for this range. Ideal sample draws do not validate a detector.

S6.4 Global access, thermometry, and resource accounting

The full states G_β, G_{β_0} are functions of one fixed Hamiltonian and commute. If every complete probe-ancilla pair is available, a common energy measurement already attains its collision value. Therefore, on two full replicas,

$$Q_2(G_\beta, G_{\beta_0}) = Q_1(G_\beta, G_{\beta_0})^2.$$

At the marked point, the full one-copy value is $2.56626736485\dots$, larger than the reduced-probe value. Full access removes the collective surplus because the complete hypotheses commute, while providing more total information. Here each access choice has its own optimized score. We have not shown that separate measurement of the full system reads the previously fixed probe observable $A_\star$ without noise. The reduced states fail to commute because the inaccessible classical ancilla label has temperature-dependent weights. This requires no entangled source and violates no data-processing inequality.

The task compares two finite temperatures through a prescribed observable. It does not estimate a continuously encoded temperature locally; single-parameter SLD Fisher information remains additive for independent probes. The coherent local thermometric response of this model was studied in Ref. [18]. We have modeled neither preparation and equilibration costs nor engine efficiency, and claim no universal strong-coupling advantage. This example has no applied work stroke or entangled input source. Giving the receiver the ancillas would change the task.

S6.5 Numerical audit and figure scope

The scripts `equilibrium_readout.py` and `verify_equilibrium_readout.py` calculate the marginal in three ways: direct global Gibbs exponentiation, the separable decomposition, and Eq. (S49). They check commutation of the global states and noncommutation of the marginals, as well as data processing and the score and moment identities. A separate calculation with `mpmath` recomputes the marked point at 70 decimal digits. Across the recorded script checks, all 1396 pass, with largest relative residual below 1.2×10^{-13} . These are implementation checks, not independent referee validation or a proof of global numerical optimality.

Main Fig. 2 scans 196 positive couplings $g/\omega \in [0.025, 1]$ at a fixed temperature pair and $\Omega/\omega = 0.2$. Each point recalculates the target and compares readout classes on that same operator. The red curve uses one feasible product implementation. We have not optimized over temperatures, Hamiltonian parameters, or all product POVMs. The plotted lines are population variances or analytic lower bounds. Ideal sample draws are stored separately. The CSV contains absolute variances and older normalized ratios; the figure uses only the absolute columns. Marked-point matrices, probabilities, and scores are in `equilibrium_marked_matrices.npz`.

S7 Data map and finite-shot scope

The two main figures come from `data/thermo_moment_bounds.csv` and `data/equilibrium_readout_scan.csv`, respectively. The script `code/make_audited_figures.py` draws both. For the defining equations and parameters, use Secs. S5–S6 rather than rounded points read from the plots. The model data come from `code/simulate_thermodynamics.py` and `code/equilibrium_readout.py`; companion scripts check the identities. Equilibrium matrices, probabilities, and scores for the marked point are in `data/equilibrium_marked_matrices.npz` and `data/equilibrium_marked.json`. No hardware data are reported.

Both plots show absolute population variance at unit response. The older normalized CSV columns do not enter either plot. At each delay or coupling, we recalculate the states and $A_\star$ before comparing readouts on that same operator. Green gives a necessary moment bound. The red curve in Fig. 2 gives

the variance of one feasible product POVM. The shading measures the arithmetic lower bound on added variance, including where the vertical axis is logarithmic.

The seeded multinomial draws in `data/monte_carlo/` sample known ideal Born probabilities, with no implemented detector. Rare outcomes and scores reaching about 1022 can make finite-sample variance estimates unstable. A sample variance below a population lower bound reflects sampling error rather than a violation of the operator inequality. Source calibration, readout error, and experimental shot budgets remain unassessed.

References

- [1] K. Temme and F. Verstraete, *J. Math. Phys.* **56**, 012202 (2015), arXiv:1112.6343 .
- [2] M. Berta, O. Fawzi, and M. Tomamichel, *Lett. Math. Phys.* **107**, 2239 (2017), arXiv:1512.02615 .
- [3] B. Regula and M. Tomamichel, Rethinking quantum smooth entropies: Tight one-shot analysis of quantum privacy amplification (2026), version 4, revised 1 September 2026, arXiv:2603.04493 .
- [4] J. M. Hammersley, *J. R. Stat. Soc. B* **12**, 192 (1950).
- [5] D. G. Chapman and H. Robbins, *Ann. Math. Stat.* **22**, 581 (1951).
- [6] Y. Tsuda and K. Matsumoto, *J. Phys. A: Math. Gen.* **38**, 1593 (2005), arXiv:quant-ph/0207150 .
- [7] K. Temme, M. J. Kastoryano, M. B. Ruskai, M. M. Wolf, and F. Verstraete, *J. Math. Phys.* **51**, 122201 (2010), corrected arXiv version 3 (2024), arXiv:1005.2358 .
- [8] A. Peres, *Phys. Rev. Lett.* **77**, 1413 (1996), arXiv:quant-ph/9604005 .
- [9] M. Horodecki, P. Horodecki, and R. Horodecki, *Phys. Lett. A* **223**, 1 (1996), arXiv:quant-ph/9605038 .
- [10] E. Chitambar, D. Leung, L. Mančinska, M. Ozols, and A. Winter, *Commun. Math. Phys.* **328**, 303 (2014), arXiv:1210.4583 .
- [11] T. Rippchen, S. Sreekumar, and M. Berta, *IEEE Trans. Inf. Theory* **71**, 6105 (2025), arXiv:2405.05037 .
- [12] H. P. Robertson, *Phys. Rev.* **34**, 163 (1929).
- [13] D. S. P. Salazar, *Phys. Rev. E* **114**, L022106 (2026), arXiv:2512.24956 .
- [14] M. Ozawa, *Annals of Physics* **311**, 350 (2004), arXiv:quant-ph/0307057 .
- [15] A. Miranowicz, M. Piani, P. Horodecki, and R. Horodecki, *Phys. Rev. A* **80**, 052303 (2009), arXiv:quant-ph/0605001 .
- [16] P. Talkner, E. Lutz, and P. Hänggi, *Phys. Rev. E* **75**, 050102 (2007), arXiv:cond-mat/0703189 .
- [17] H. J. D. Miller and J. Anders, *Nat. Commun.* **9**, 2203 (2018), arXiv:1801.08057 .
- [18] A. Ullah, M. T. Naseem, and Ö. E. Müstecaplıoğlu, *Phys. Rev. Research* **5**, 043184 (2023), arXiv:2211.05461 .
- [19] G. M. D’Ariano and P. Perinotti, *Phys. Rev. Lett.* **98**, 020403 (2007), arXiv:quant-ph/0610058 .